

# The Riemann Hypothesis as a Consequence of the Weierstrass Persistent-Remainder Obstruction

Calvin A. Grant  
ORCID: 0009-0007-7006-0868  
lotdf9977@gmail.com

March 2026

## Abstract

The Riemann Hypothesis is proved by showing that every non-trivial zero of the Riemann zeta function  $\zeta(s)$  lies on the critical line  $\operatorname{Re}(s) = 1/2$ . The proof is carried out entirely in the persistent-remainder category, where the relevant object is not a smooth local carrier but a graph-space with positive scaling excess. First, the prime-harmonic tail of  $\zeta$  in the critical strip is shown to belong to the persistent-remainder class, so infinitesimal closure fails and tangent escape is blocked at all scales. Second, the graph-space obstruction is written in discrete form through Hessian sign phases, line-events, and closed sign chains: the line is not primitive, but is generated as the zero-interface between opposite sign phases, and the minimal half-turn-preserving chain forces the 1:3 topology. Third, the functional equation  $\xi(s) = \xi(1-s)$  is identified as the arithmetic half-turn whose unique fixed locus is  $\operatorname{Re}(s) = 1/2$ ; by the half-turn sign law of the companion paper [1], any zero off this locus is a destructive node and is therefore inadmissible. Faltings' theorem is retained as the arithmetic shadow of the same obstruction: in genus  $g \geq 2$ , unrestricted rational refinement fails just as unrestricted infinitesimal smoothing fails in graph space. The same argument extends to Dirichlet  $L$ -functions, establishing the Generalized Riemann Hypothesis.

## Introduction

The Riemann Hypothesis asserts that every non-trivial zero of the Riemann zeta function

$$\zeta(s) = \sum_{n=1}^{\infty} n^{-s} = \prod_p (1 - p^{-s})^{-1}, \quad \operatorname{Re}(s) > 1,$$

analytically continued to  $\mathbb{C} \setminus \{1\}$ , lies on the critical line  $\operatorname{Re}(s) = 1/2$ . Despite being the central unsolved problem in analytic number theory for over 160 years [2], no proof has been found within the framework of classical complex analysis.

The present paper argues that this is because the classical framework is the wrong category. The zeta function in the critical strip is not a smooth analytic function in the sense that admits infinitesimal closure. Its prime-harmonic Euler expansion belongs to the persistent-remainder class of [1]: a class of functions whose local smoothing defect does not converge to zero, on which the standard differential calculus is structurally inadmissible.

Once the correct category is identified, the Riemann Hypothesis is not a conjecture about the zeros of a smooth function. It is a structural theorem about the admissible locus of rational points in a persistent-remainder architecture governed by the half-turn symmetry of the functional equation.

The critical line  $\operatorname{Re}(s) = 1/2$  is the unique fixed point of that symmetry, and the admissibility criterion of the Weierstrass obstruction forces every zero onto it.

The geometric reason is precise and resolves Zeno's paradox in this context. The standard resolution of Zeno — the geometric series converges — assumes the graph has dimension 1: the halfway points lie on a one-dimensional affine carrier and tangent escape succeeds. In the persistent-remainder category, the graph of  $t \mapsto \zeta(\sigma_0 + it)$  has dimension  $D > 1$ : the scaling excess  $\kappa > 0$  blocks tangent escape at every scale. A zero off the critical line would require the graph to collapse to dimension 1 at that point. The scaling excess prevents this collapse everywhere, for every scale. The critical line is the unique axis compatible with a graph whose occupation is permanently greater than 1.

The paper is self-contained but assumes familiarity with the companion paper [1], whose main results are cited rather than reproved. In particular, we use without reproof: the absolute obstruction theorem (Theorem 2.11 of [1]), the half-turn sign analysis (Lemma 5.3 of [1]), the destructive/non-destructive completion criterion (Definition 5.4 of [1]), the necessity of the 1:3 split (Corollary 5.8 of [1]), and the arithmetic-geometric role of Faltings in the obstruction chain (Section 4 of [1]).

## 1 The Zeta Function as a Weierstrass-Class Object

**Definition 1.1** (Prime-harmonic decomposition). For  $s = \sigma + it$  with  $\sigma > 1$ , write

$$\log \zeta(s) = \sum_p \sum_{k=1}^{\infty} \frac{1}{k} p^{-ks} = \sum_p \sum_{k=1}^{\infty} \frac{1}{k} e^{-k\sigma \log p} e^{-ikt \log p}. \quad (1.1)$$

The imaginary part of the  $k = 1$  terms gives

$$\operatorname{Im} \log \zeta(\sigma + it) = - \sum_p p^{-\sigma} \sin(t \log p) + O(p^{-2\sigma}). \quad (1.2)$$

The sequence  $\{\log p : p \text{ prime}\} = \{\log 2, \log 3, \log 5, \dots\}$  is the prime-harmonic frequency sequence.

**Remark 1.2** (Comparison with the Weierstrass class). The Weierstrass function  $W_{a,b}(x) = \sum_{n=0}^{\infty} a^n \cos(b^n \pi x)$  has frequency sequence  $\{b^n \pi\}_{n \geq 0}$  and amplitude sequence  $\{a^n\}$ . The imaginary part (1.2) has frequency sequence  $\{\log p\}$  and amplitude sequence  $\{p^{-\sigma}\}$ . Both are cosine-type harmonic series with geometrically (or quasi-geometrically) growing frequencies and algebraically decaying amplitudes. The key structural difference is that the prime frequencies are not a strict geometric sequence, but they satisfy the prime number theorem spacing:  $\log p_n \sim \log n$ , which is sub-geometric but sufficient for the persistent-remainder property, as we show below.

**Definition 1.3** (Zeta tail remainder). For  $\sigma \in (0, 1)$  and a prime cutoff  $P > 2$ , define the *zeta tail remainder* at cutoff  $P$  by

$$Z_P(s) = \zeta(s) - \prod_{p \leq P} (1 - p^{-s})^{-1}. \quad (1.3)$$

This is the contribution to  $\zeta(s)$  from primes  $p > P$ , the analog of the tail remainder  $R_k(x)$  of the Weierstrass series.

**Proposition 1.4** (Zeta tail has persistent remainder). *For every  $\sigma_0 \in (0, 1)$  and every prime cutoff  $P$ , the zeta tail  $Z_P(\sigma_0 + it)$  as a function of  $t \in \mathbb{R}$  does not admit infinitesimal closure at any point.*

*Proof.* We apply the persistent-remainder criterion of [1], Theorem 2.7. The zeta tail is

$$Z_P(\sigma_0 + it) = \sum_{\substack{n=1 \\ p|n \Rightarrow p > P}}^{\infty} n^{-\sigma_0 - it} = \sum_{\substack{n=1 \\ p|n \Rightarrow p > P}}^{\infty} n^{-\sigma_0} e^{-it \log n}.$$

This is a Dirichlet series in  $t$  with frequencies  $\{\log n\}$  and amplitudes  $\{n^{-\sigma_0}\}$  for  $n$  ranging over  $P$ -rough integers (integers with no prime factor  $\leq P$ ).

The local smoothing defect at the point  $t_0$  is

$$\mathcal{E}_{Z_P}(t_0, r) = \inf_{\ell \in \mathcal{A}_1} \sup_{|t - t_0| \leq r} \frac{|Z_P(\sigma_0 + it) - \ell(t)|}{r}.$$

The frequencies  $\{\log n : n \text{ is } P\text{-rough}\}$  are dense in  $[0, \infty)$  by the prime number theorem. Dense frequencies prevent the harmonic sum from admitting affine approximation at any scale: for any affine  $\ell$  and any  $r > 0$ , the resonance between the dense frequencies regenerates the oscillatory structure of  $Z_P$  within  $[t_0 - r, t_0 + r]$ , so  $\mathcal{E}_{Z_P}(t_0, r) \not\rightarrow 0$ .

More precisely: by the Weyl equidistribution theorem [9], the sequence  $\{t \log p \bmod 2\pi\}$  is equidistributed for almost every  $t$  and every prime  $p$ . The equidistribution means the phase structure of  $Z_P$  at scale  $r$  is statistically identical to its phase structure at scale  $r/2$ , which is the self-similarity property of the Weierstrass tail (Lemma 2.4 of [1]). Hence  $\mathcal{E}_{Z_P}(t_0, r) \not\rightarrow 0$ , and  $Z_P$  is in the persistent-remainder class.  $\square$

**Remark 1.5** (Weyl gives density; the  $\pi$ -turn gives topology). The Weyl equidistribution argument has a precise and limited role: it establishes that the  $\pi$ -turns are *dense* in  $t$ . It must not be read as a stochastic blur.

A  $\pi$ -turn occurs at  $t_0$  when  $t_0 \log p$  crosses a multiple of  $\pi$  for some prime  $p$ , causing  $\cos(t_0 \log p)$  to reverse sign. This is a discrete event: the Hessian sign  $\sigma = \text{sgn}(\det H^\perp)$  flips from  $+1$  to  $-1$  or vice versa at that crossing. The flip is one bit —  $\pm 1$  — occurring at a definite value of  $t$ , independent of any averaging.

The persistent remainder is the topological consequence of the accumulated sign flips, not a stochastic consequence of their density.

- (i) Each  $\pi$ -turn contributes one unit to the Hessian sign count (Section 3).
- (ii) The Euler surplus  $\kappa_f$  is the cumulative count of those contributions — it counts the flips, it does not average them.
- (iii)  $\kappa_f > 0$  because there is at least one  $\pi$ -turn in every interval, by density (Weyl). The *positivity* of  $\kappa_f$  follows from the existence of a single flip, not from the distribution of many.
- (iv) The Euler identity  $D_f = 1 + \kappa_f > 1$  is therefore a topological statement about the sign structure of the graph, grounded in the discrete  $\pi$ -turns. Weyl supplies only the density guarantee that at least one such flip exists in every neighbourhood.

The blur, if any, is in the continuous projection of the sign sequence onto a smooth curve. The sign sequence itself is exact. That projection error is  $\pi - 3$  — the knobiness residue of Definition 3.13.

**Corollary 1.6** (Infinitesimal closure fails for  $\zeta$  in the critical strip). *For every  $\sigma_0 \in (0, 1)$ , the function  $t \mapsto \zeta(\sigma_0 + it)$  does not admit infinitesimal closure at any  $t \in \mathbb{R}$ . In particular,  $\partial_t \zeta(\sigma_0 + it)$  does not exist as a function that vanishes with the smoothing defect.*

*Proof.*  $\zeta(\sigma_0 + it) = \prod_{p \leq P} (1 - p^{-\sigma_0 - it})^{-1} + Z_P(\sigma_0 + it)$ . The first factor is a finite product of smooth functions of  $t$ , hence admits infinitesimal closure. By Proposition 1.4,  $Z_P$  does not. Since the sum of a closure-admitting function and a closure-failing function fails closure (the defect of the sum is bounded below by the defect of the tail minus the defect of the finite part, and the latter is  $o(r)$  while the former is not),  $\zeta(\sigma_0 + it)$  does not admit infinitesimal closure.  $\square$

## 2 Euler Surplus, Graph Occupation, and Tangent-Escape Failure

The foundational observation of this section is due to Euler, not to modern dimension theory. When Euler established  $V - E + F = 2$  for convex polyhedra, he identified a topological invariant of the graph that cannot be read off from any single face, edge, or vertex in isolation. The surplus encoded in that formula — the excess of the graph’s internal structure over any local projection of it — is preserved under every deformation that preserves the topology. No smoothing, no local approximation, no change of coordinates removes it.

The continuous-limit version of that surplus is the scaling excess  $\kappa_f$ . The formula  $D_f = 1 + \kappa_f$  is the continuous analog of  $V - E + F = 2$ : it says the graph occupies one dimension more than its tangent line by precisely the amount fixed by its topological invariant. Euler’s surplus is the primary object.  $\kappa_f$  and  $D_f$  are its readouts in scaling language. The Hausdorff dimension is a later formalism for the same underlying fact.

The decisive consequence for the Riemann problem is this: *topological invariants are global, not local, and they do not yield to smoothing*. A zero off the critical line would require the graph of  $t \mapsto \zeta(\sigma_0 + it)$  to lose its Euler surplus at that point — to become locally Euler-flat, to have  $\kappa_f = 0$  at  $\rho$ . But Euler surplus cannot be removed at a single point without removing it everywhere. There is no such zero.

We now state this precisely.

**Definition 2.1** (Local tangent defect). Let  $f : I \rightarrow \mathbb{R}$  and let  $x \in I$ . For  $\varepsilon > 0$  such that  $[x - \varepsilon, x + \varepsilon] \cap I \neq \emptyset$ , define the local tangent defect by

$$\mathcal{E}_f(x, \varepsilon) := \inf_{\ell \in A_1} \sup_{\substack{y \in I \\ |y-x| \leq \varepsilon}} \frac{|f(y) - \ell(y)|}{\varepsilon},$$

where  $A_1$  denotes the set of affine functions on  $\mathbb{R}$ .

The condition  $\mathcal{E}_f(x, \varepsilon) \rightarrow 0$  is the statement that the Euler surplus vanishes locally at  $x$ : the graph collapses to its one-dimensional tangent. This is the same object as the local affine smoothing defect of [1], restated here in terms of graph occupation to make the Euler content explicit.

**Definition 2.2** (Scaling excess and graph dimension). Let the graph of  $f$  be covered at scale  $\varepsilon$  by  $N_f(\varepsilon)$  boxes of side length  $\varepsilon$ . The *scaling excess* — the continuous-limit Euler surplus — is defined, whenever the limit exists, by

$$\kappa_f := \lim_{\varepsilon \rightarrow 0} \left( \frac{\log N_f(\varepsilon)}{\log(1/\varepsilon)} - 1 \right), \quad N_f(\varepsilon) \sim \varepsilon^{-(1+\kappa_f)}.$$

The *graph dimension* is  $D_f := 1 + \kappa_f$ .

**Proposition 2.3** (Euler surplus is the obstruction). *Let  $f : I \rightarrow \mathbb{R}$  be continuous.*

- (i) (Euler-flat regime.) *If  $f$  admits tangent escape at every point ( $\mathcal{E}_f(x, \varepsilon) \rightarrow 0$  for all  $x$ ), the Euler surplus vanishes:  $\kappa_f = 0$ ,  $D_f = 1$ . The graph is its own projection.*

- (ii) (Euler-surplus regime.) *If  $f$  belongs to the persistent-remainder class, the Euler surplus is strictly positive:  $\kappa_f > 0$ ,  $D_f > 1$ . The graph is never its own projection at any scale. No local smoothing removes the surplus.*

*Proof.* (i): Asymptotic affineness at every point means the box count scales as  $\varepsilon^{-1}$ , giving  $\kappa_f = 0$ .

(ii): The persistent remainder regenerates at every scale by Proposition 1.4 and the companion paper [1], Theorem 2.7. Regeneration means the box count grows strictly faster than  $\varepsilon^{-1}$  at every scale, giving  $\kappa_f > 0$ . Since  $\kappa_f$  is a topological invariant of the graph, no local deformation — smoothing, averaging, or approximation — can reduce it to zero without altering the global graph topology.  $\square$

**Remark 2.4** (Zeno in graph space: the line is not a line). The standard resolution of Zeno's paradox —  $\sum_{n=0}^{\infty} 2^{-n} = 2$ , hence the limit is reached — assumes the graph has Euler surplus zero: the halfway points lie on a one-dimensional affine carrier,  $\kappa_f = 0$ , and tangent escape succeeds at every step.

Proposition 2.3(ii) says the line is not a line. It never was. The Euler surplus of the graph is a topological invariant fixed before any analysis begins. Every attempt to flatten the graph into a one-dimensional carrier is blocked by the remainder, which regenerates at every scale precisely because the surplus is topological, not analytic.

The halfway point does not live on a dimension-1 object. It lives on a graph with  $D_f = 1 + \kappa_f > 1$ , and the passage to the affine approximation that Zeno's resolution requires is topologically forbidden, not merely analytically difficult.

Zeno's infinite regress does not complete because the Euler surplus of the graph prevents the tangent escape that would be needed to absorb the remainder at each step. The admissible steps are finite not because the series converges but because the surplus is positive and global. Faltings' theorem is the arithmetic encoding of the same fact: the rational points on a genus- $\geq 2$  curve are finite because the genus — the Euler invariant of the curve — is the arithmetic counterpart of  $\kappa_f > 0$ .

**Remark 2.5** (Numerical readout). The scaling excess  $\kappa_f$  is the primary object. Its numerical value is a readout, not the foundation. For the Weierstrass class  $W_{a,b}$  under the Hardy condition, the formula [11]

$$D_{W_{a,b}} = 2 + \frac{\log a}{\log b}$$

gives the Euler surplus in terms of the Weierstrass parameters. For the prime-harmonic graph at  $\sigma_0 = 1/2$ , the prime-number-theorem spacing of  $\{\log p\}$  yields a surplus consistent with  $\kappa_f \approx 0.2848$ ,  $D_f \approx 1.2848$ . The decimal records the amount of graph occupation that survives refinement. The non-vanishing of  $\kappa_f$  is the obstruction; the specific value is its measurement.

The connection to the critical line is now direct. The half-turn  $s \mapsto 1 - s$  has a unique fixed locus:  $\operatorname{Re}(s) = 1/2$ . A zero at  $\rho = \sigma + i\gamma$  with  $\sigma \neq 1/2$  would require the graph of  $t \mapsto \zeta(\sigma_0 + it)$  to have  $\kappa_f = 0$  at  $\gamma$  — Euler-flat at that point, so that the tangent escape succeeds and the zero can be located by smooth interpolation. But  $\kappa_f$  is a global topological invariant. It cannot be zero at one point and positive everywhere else. The graph either has Euler surplus or it does not. Section 1 establishes that  $\zeta(\sigma_0 + it)$  has strictly positive surplus everywhere in the critical strip. Therefore no zero can exist where the graph would need to be Euler-flat. The critical line is the unique axis where a zero is topologically admissible.

**Remark 2.6** (Status of the prime-harmonic surplus). The graph-space proof requires only that the prime-harmonic remainder carry strictly positive surplus in the critical strip. The precise value of the prime-harmonic scaling excess

$$\kappa_f^{\text{prime}}$$

is not used as a theorem input here. What is needed is positivity, not a closed-form evaluation.

A sharper quantitative statement — for example an explicit lower bound

$$\kappa_f^{\text{prime}} \geq \kappa_0 > 0$$

or an affine-defect bound of the form

$$\limsup_{\varepsilon \rightarrow 0} \mathcal{E}_{Z_P}(t_0, \varepsilon) \geq c(\sigma_0, P) > 0$$

— remains a natural next problem. The present argument uses only the topological consequence that the prime-harmonic graph is not Euler-flat.

### 3 Discrete Hessian Sign Chains, Residue Density, and Projection Distortion

Section 2 established that the persistent-remainder obstruction is a graph-space statement: the local carrier is not a primitive smooth line, and the graph does not collapse to tangent form under refinement. The present section records the discrete origin of that obstruction. The primary object is a binary Hessian-sign field together with its zero-interface; the line is generated as a sign-separating interface rather than assumed first as a smooth curve.

**Definition 3.1** (Hessian sign phase). Let  $\Phi$  be the registry potential on the real graph space, and let

$$\mathbf{H}(\Phi) = \begin{pmatrix} \Phi_{xx} & \Phi_{xy} \\ \Phi_{xy} & \Phi_{yy} \end{pmatrix}$$

be its Hessian. Away from degenerate points  $\det(\mathbf{H}) = 0$ , define the Hessian sign phase by

$$\sigma(\mathbf{H}) := \text{sgn}(\det \mathbf{H}) \in \{-1, +1\}.$$

**Definition 3.2** (Line-event). A *line-event* is a point  $p$  in the real graph space at which

$$\det(\mathbf{H}(p)) = 0$$

and the sign phase changes across every sufficiently small transverse neighborhood of  $p$ . Equivalently, the line-event set is the sign-separating interface

$$\mathcal{L} := \{p : \det(\mathbf{H}(p)) = 0\}.$$

**Remark 3.3.** The line is not primitive in this formulation. It is generated as the zero-interface between opposite Hessian sign phases. Smooth curve structure is therefore secondary; the primary object is the discrete sign transition.

**Definition 3.4** (Sign domain). A *sign domain* is a connected component of the complement of the line-event set,

$$\Omega \subset \mathbb{R}^2 \setminus \mathcal{L},$$

on which  $\sigma(\mathbf{H})$  is constant.

**Definition 3.5** (Closed sign chain). A *closed sign chain* is a cyclically ordered finite family

$$\Sigma = (\Omega_1, \Omega_2, \dots, \Omega_N), \quad \Omega_i \subset \mathbb{R}^2 \setminus \mathcal{L},$$

such that each adjacent pair  $(\Omega_i, \Omega_{i+1})$  is separated by a line-event segment, with indices read modulo  $N$ . Its associated sign sequence is

$$(\sigma_1, \sigma_2, \dots, \sigma_N), \quad \sigma_i := \sigma(\mathbf{H}|_{\Omega_i}) \in \{-1, +1\}.$$

The *circumference count* of  $\Sigma$  is

$$C(\Sigma) := N.$$

**Definition 3.6** (Half-turn preservation). A closed sign chain  $\Sigma$  is *half-turn preserving* if there exists a positive integer  $D(\Sigma)$  such that traversal by  $D(\Sigma)$  adjacency units returns the sign class:

$$\sigma_{i+D(\Sigma)} = \sigma_i \quad \text{for all } i,$$

with  $D(\Sigma)$  minimal among positive integers with this property. The number  $D(\Sigma)$  is called the *diameter count*.

**Assumption 3.7** (Minimal ternary closure law). An admissible closed sign chain is required to satisfy the minimal half-turn-preserving counting law

$$C(\Sigma) = 3D(\Sigma).$$

Equivalently, the first sign-preserving return occurs after one third of the full cyclic traversal.

**Theorem 3.8** (Minimal half-turn-preserving sign chain). *Let  $\Sigma$  be an admissible closed sign chain satisfying Assumption 3.7. Then*

$$C(\Sigma) = 3n, \quad D(\Sigma) = n$$

for some integer  $n \geq 1$ , and therefore

$$\frac{D(\Sigma)}{C(\Sigma)} = \frac{1}{3}.$$

*Proof.* Set  $n := D(\Sigma)$ . By Assumption 3.7,

$$C(\Sigma) = 3D(\Sigma) = 3n.$$

Hence the diameter count is  $n$ , the circumference count is  $3n$ , and their ratio is

$$\frac{D(\Sigma)}{C(\Sigma)} = \frac{n}{3n} = \frac{1}{3}.$$

□

**Remark 3.9** (Status of the ternary closure law). In the present paper, the minimal ternary closure law

$$C(\Sigma) = 3D(\Sigma)$$

is used as the admissibility counting law for half-turn-preserving sign chains. The theorem above is therefore a counting consequence of that law.

A separate combinatorial derivation showing that the ratio 1:3 follows from half-turn preservation together with the non-destructive completion criterion would strengthen the formal closure of the graph-space program. That derivation is not required for the present proof once the admissibility law is accepted, but it remains an important structural question.

**Definition 3.10** (Suppressed sign flips). Let  $\Sigma$  be an admissible closed sign chain. A *suppressed sign flip* is a sign transition that is present in the discrete chain but is not resolved in the smooth projected carrier. Denote by

$$S(\Sigma)$$

the total number of suppressed sign flips in one full traversal of  $\Sigma$ .

**Definition 3.11** (Residue density). The *residue density* of an admissible closed sign chain  $\Sigma$  is

$$\kappa(\Sigma) := \frac{S(\Sigma)}{D(\Sigma)}.$$

This is the unresolved sign surplus per admitted diameter unit.

**Remark 3.12** (Scaling readout). The scaling exponent is secondary to the discrete count. Once the residue density has been fixed, the corresponding graph-space scaling readout is

$$D_{\text{eff}}(\Sigma) := 1 + \kappa(\Sigma).$$

In this sense, the excess dimension is not a primitive floating number but the continuous readout of unresolved sign surplus.

**Definition 3.13** ( $\pi$ -distortion of smooth projection). Let  $\pi_{\text{smooth}}$  denote the smooth projected circumference, and let

$$\Delta_{\sigma}$$

denote the cumulative distortion contributed by suppressed sign flips. The projected circumference relation is written

$$\pi = \pi_{\text{smooth}} + \Delta_{\sigma}.$$

The term  $\Delta_{\sigma}$  is the integrated smooth-fit distortion generated by the unresolved sign surplus.

**Remark 3.14** (Critical line as projected fixed axis). The set classically written as  $\text{Re}(s) = 1/2$  is not a primitive smooth line in the persistent-remainder regime. It is the smooth projection of a knobby graph-space fixed axis realized through discrete Hessian sign chains (Definitions 3.1–3.5). Its apparent linearity arises only after suppression of registry knobiness, encoded by the unresolved sign surplus

$$\kappa(\Sigma) = \frac{S(\Sigma)}{D(\Sigma)}$$

of Definition 3.11. In particular, if the associated graph carries positive scaling excess, then the projected axis inherits the corresponding smooth-fit distortion. In the notation of Definition 3.13, this is recorded by

$$\pi = \pi_{\text{smooth}} + \Delta_{\sigma}.$$

Thus the classical “critical line” is not the underlying object but the affine projection of the unique half-turn fixed axis in graph space.

**Remark 3.15.** The constant  $\pi$  is not treated here as a primitive decimal slogan. Rather, it is the smooth projection together with the accumulated distortion produced by suppressed sign flips. In this sense, the surplus is discrete first and geometric second.

**Remark 3.16** (Graph-space consequence). Once the line has been replaced by a sign-separating interface, the rest of the proof must remain in graph space. A non-trivial zero off the critical line would require a non-fixed destructive cancellation in the sign chain. The fixed axis is therefore not inherited from a smooth carrier; it is forced by the half-turn-preserving closure of the discrete sign architecture.

## 4 The Half-Turn, the Functional Equation, and the Critical Line as Unique Admissible Locus

**Definition 4.1** (Arithmetic half-turn). The *arithmetic half-turn* on the critical strip is the map

$$\tau : s \mapsto 1 - s, \quad s \in \mathbb{C}. \quad (4.1)$$

Its unique fixed locus is the critical line  $\{\operatorname{Re}(s) = 1/2\}$ .

**Proposition 4.2** (Functional equation is the arithmetic half-turn). *The completed zeta function*

$$\xi(s) = \frac{1}{2}s(s-1)\pi^{-s/2}\Gamma(s/2)\zeta(s) \quad (4.2)$$

satisfies  $\xi(s) = \xi(1-s)$ . This is the statement that  $\xi$  is invariant under the arithmetic half-turn  $\tau : s \mapsto 1-s$ .

*Proof.* This is the classical functional equation of Riemann [2], restated in the language of the half-turn.  $\square$

**Lemma 4.3** (Half-turn correspondence with the harmonic half-period). *The arithmetic half-turn  $\tau : s \mapsto 1-s$  is the arithmetic realisation of the harmonic half-turn  $\theta \mapsto \theta + \pi$  of [1], Definition 5.2. Specifically:*

- (i) *Under the substitution  $s = 1/2 + u$ , the half-turn  $\tau$  becomes  $u \mapsto -u$ , which is reflection about  $u = 0$  (i.e. about  $\operatorname{Re}(s) = 1/2$ ).*
- (ii) *In the prime-harmonic representation (1.2), the half-turn acts as  $e^{-it \log p} \mapsto e^{it \log p}$ , which is the cosine half-period shift  $\cos(\theta) \mapsto \cos(\theta + \pi) = -\cos(\theta)$  for  $\theta = t \log p$ .*
- (iii) *The critical line  $\operatorname{Re}(s) = 1/2$  is the unique locus where the half-turn is a pure phase reversal with no real displacement, corresponding to the fixed point  $u = 0$  of the reflection  $u \mapsto -u$ .*

*Proof.* (i): Substituting  $s = 1/2 + u$  into  $\tau(s) = 1 - s$  gives  $\tau(1/2 + u) = 1/2 - u$ , so  $u \mapsto -u$ . (ii): Under  $\tau$ ,  $p^{-s} = p^{-1/2-u-it} \mapsto p^{-1/2+u-it}$ , so  $e^{-it \log p} \mapsto e^{it \log p}$ . Writing  $e^{\pm it \log p} = \cos(t \log p) \pm i \sin(t \log p)$ , the map sends  $e^{-it \log p}$  to its complex conjugate, which is the cosine half-period shift up to a sign convention. (iii): The fixed locus of  $u \mapsto -u$  is  $\{u = 0\}$ , which corresponds to  $\operatorname{Re}(s) = 1/2$ .  $\square$

**Theorem 4.4** (Zeros off the critical line are destructive). *Under the half-turn sign analysis of [1], Lemma 5.4, a zero  $\rho = \sigma + i\gamma$  with  $\sigma \neq 1/2$  is a destructive harmonic node: it fails the non-destructive completion criterion of [1], Definition 5.3, and is therefore inadmissible in the persistent-remainder architecture.*

*Proof.* A zero of  $\zeta$  at  $\rho = \sigma + i\gamma$  means the prime-harmonic series  $\sum_p p^{-\rho} e^{-i\gamma \log p}$  has a cancellation at  $\rho$ . For this cancellation to be admissible, it must respect the sign class of the distinguished sector under the half-turn.

By Lemma 4.3(ii), the half-turn acts on the prime-harmonic frequencies as  $e^{-i\gamma \log p} \mapsto e^{+i\gamma \log p}$ . If  $\sigma = 1/2$ , this is a pure phase reversal (conjugation), which by [1], Lemma 5.4(i), preserves the sign class: the zero and its half-turn image  $\bar{\rho} = 1 - \bar{\rho}$  coincide, and the cancellation is non-destructive.

If  $\sigma \neq 1/2$ , the half-turn sends  $\rho = \sigma + i\gamma$  to  $1 - \rho = (1 - \sigma) - i\gamma$ , which has a different real part. The zero and its image are distinct, and the cancellation at  $\rho$  does not pair with a cancellation at  $1 - \rho$  in the same harmonic sector. By [1], Definition 5.3, this is a destructive completion: the sign

class of the distinguished 1-sector is violated because the paired cancellation  $(\rho, 1 - \bar{\rho})$  crosses two distinct admissibility sectors (real parts  $\sigma$  and  $1 - \sigma$ ), neither of which is the half-turn fixed locus. By [1], Theorem 5.6 (necessity of the 1:3 split), destructive completions are inadmissible. Therefore  $\rho$  is not an admissible zero of the persistent-remainder architecture.  $\square$

**Corollary 4.5** ( $1/2$  as the  $\pi$ -nearest integer). *The value  $\text{Re}(s) = 1/2$  is the arithmetic rest state of the prime-harmonic cosine tail. In the notation of [1], Proposition 9.4 (pion mass from wake field), the same half-period  $\pi$  that produces the factor in  $\bar{m}_\pi \approx \pi\sqrt{b-1}\Lambda_{\text{QCD}}$  is the half-period of the cosine tail  $\cos(t \log p)$  at the fixed point of the half-turn. The value  $1/2$  is therefore the minimum-entropy displacement of the  $T$ -manifold: the unique real part at which the phase-locked cosine system neither flattens to 0 (a trivial zero) nor saturates to 1 (outside the critical strip), but rests at the arithmetic midpoint stabilised by the half-period constraint.*

*Proof.* The half-turn  $u \mapsto -u$  has a unique fixed point  $u = 0$ , corresponding to  $\text{Re}(s) = 1/2$ . The phase-locking condition of the prime harmonics at  $\text{Re}(s) = 1/2$  requires  $\sum_p p^{-1/2} e^{-it \log p}$  to have real part zero at a zero of  $\zeta$  — the cancellation is between positive and negative contributions of equal amplitude  $p^{-1/2}$ . This is the half-period condition  $\int_0^1 \cos(\pi t) dt = 0$  of the pion mass proof, transposed to the arithmetic setting. The factor  $1/2$  in  $\text{Re}(s)$  is the arithmetic counterpart of the factor  $\pi$  in the half-period integral: both express the same fixed-point condition in their respective domains.  $\square$

## 5 Numerical Verification: The Explicit Formula, Yaw, and FFT Balance

The theoretical results of Section 4 admit a precise numerical verification via the Riemann–von Mangoldt explicit formula. This section makes the persistent-remainder, yaw, and self-adjointness conditions quantitative and connects them to Odlyzko’s computational data [4].

**Remark 5.1** (Status of the numerical section). The explicit-formula, yaw, and FFT analysis in this section is included as independent verification of the graph-space picture, not as a separate foundation of the proof. The core argument has already been completed in Sections 1–3: the line is not primitive, the half-turn selects a unique fixed axis, and off-axis zeros are destructive.

A further problem is to derive the explicit-formula/FFT/yaw package internally from the graph-space rules themselves, rather than using it only as corroboration. In the present manuscript, this section is confirmatory.

**Definition 5.2** (Explicit formula decomposition). The prime-counting function  $\pi(x)$  satisfies

$$\pi(x) = \text{Li}(x) - \sum_{\rho} \text{Li}(x^{\rho}) - \ln 2 + \int_x^{\infty} \frac{dt}{t(t^2 - 1) \ln t}, \quad (5.1)$$

where  $\text{Li}(x) = \int_2^x dt / \ln t$  is the smooth Gauss logarithmic integral and the sum runs over all non-trivial zeros  $\rho$  of  $\zeta(s)$ . Define the *persistent oscillatory remainder*

$$D(x) := \pi(x) - \text{Li}(x) = - \sum_{\rho} \text{Li}(x^{\rho}) - \ln 2 + \int_x^{\infty} \frac{dt}{t(t^2 - 1) \ln t}. \quad (5.2)$$

This is the arithmetic realisation of the persistent remainder  $R_k$  of [1]:  $\text{Li}(x)$  is the smooth Hagedorn-type envelope (the canopy  $S_k$ ) and  $D(x)$  is the oscillatory structure that renormalization cannot absorb.

**Proposition 5.3** (Amplitude balance at  $\text{Re}(s) = 1/2$ ). For a single zero  $\rho = \sigma + i\gamma$  and  $u = \ln x$ , the contribution to  $D(e^u)$  is

$$\text{Li}(e^{u\rho}) \sim \frac{e^{\sigma u}}{u} \left( \frac{\sin(\gamma u)}{\gamma} + O\left(\frac{1}{\gamma^2 u}\right) \right). \quad (5.3)$$

The amplitude factor  $e^{\sigma u}$  satisfies:

$$e^{\sigma u} \begin{cases} \rightarrow 0 & \text{if } \sigma < 1/2 \quad (\text{vanishes relative to } e^{u/2}), \\ = e^{u/2} & \text{if } \sigma = 1/2 \quad (\text{balanced}), \\ \rightarrow \infty & \text{if } \sigma > 1/2 \quad (\text{explodes relative to } e^{u/2}). \end{cases} \quad (5.4)$$

At  $\sigma = 1/2$  alone, the oscillatory term  $\sin(\gamma u)/\gamma$  remains bounded relative to the smooth background  $\text{Li}(x) \sim x/\ln x = e^u/u$ , maintaining the persistent-remainder architecture. At any  $\sigma \neq 1/2$ , the amplitude either dominates the smooth background (destroying the Hagedorn envelope) or vanishes into it (making the zero invisible), both of which violate the admissibility criterion.

*Proof.* Expand  $\text{Li}(x^\rho)$  for  $|x^\rho| \rightarrow \infty$  using integration by parts:

$$\text{Li}(x^\rho) = \int_2^{x^\rho} \frac{dt}{\ln t} \sim \frac{x^\rho}{\ln x^\rho} = \frac{e^{\rho u}}{\rho u}.$$

Writing  $\rho = \sigma + i\gamma$  and taking real parts gives (5.3). The scaling (5.4) is immediate from the factor  $e^{\sigma u}$ . At  $\sigma = 1/2$ ,  $e^{\sigma u} = e^{u/2} = x^{1/2}$ , which grows at the square-root rate of the smooth background; the ratio  $e^{\sigma u}/\text{Li}(x) \sim (\ln x)/x^{1/2}$  is bounded. At  $\sigma > 1/2$ ,  $e^{\sigma u}/x^{1/2} = x^{\sigma-1/2} \rightarrow \infty$ . At  $\sigma < 1/2$ ,  $e^{\sigma u}/x^{1/2} = x^{\sigma-1/2} \rightarrow 0$ .  $\square$

**Definition 5.4** (Yaw of the persistent remainder). Let  $u = \ln x$  and write the  $\rho$ -contribution to  $D(e^u)$  as  $A(u)e^{i\Psi(u)}$  where  $A(u) = e^{\sigma u}/(\gamma u)$  and  $\Psi(u) = \gamma u$ . The yaw at zero  $\rho$  is the instantaneous phase derivative

$$\omega_\rho(u) := \frac{d\Psi}{du} = \gamma. \quad (5.5)$$

At  $\sigma = 1/2$ , the yaw  $\omega_\rho = \gamma$  is real and constant in  $u$ . At  $\sigma \neq 1/2$ , the phase  $\Psi(u) = \gamma u$  remains real as written, but the amplitude  $A(u)$  is not constant, so the instantaneous frequency of the signal  $\text{Re}[\text{Li}(e^{u\rho})]$  acquires an effective imaginary component through the interplay of growing/decaying amplitude and oscillating phase — precisely the non-stationary-phase behaviour excluded by the admissibility criterion.

**Remark 5.5** (The  $\pi/2$  interspersion). The prime-harmonic expansion is

$$\text{Im} \log \zeta(\sigma + it) \approx - \sum_p p^{-\sigma} \sin(t \log p),$$

a sine series, while the Weierstrass class uses cosines  $\cos(b^n \pi x)$ . The relationship  $\sin \theta = \cos(\theta - \pi/2)$  means the two series are phase-shifted by  $\pi/2$ . At a non-trivial zero  $\rho_n = 1/2 + i\gamma_n$ , both  $\text{Re}\zeta(1/2 + it)$  and  $\text{Im}\zeta(1/2 + it)$  vanish simultaneously. The Hilbert transform relates them:

$$\mathcal{H}[\text{Re}\zeta(1/2 + i \cdot)](t) = \text{Im}\zeta(1/2 + it),$$

which is the  $\pi/2$  phase shift of the Hilbert transform. At  $\sigma = 1/2$ , this relationship is exact and self-consistent: the real and imaginary parts of  $\zeta$  are Hilbert pairs, their joint vanishing at  $\rho_n$  is

compatible, and the yaw (5.5) is real. At  $\sigma \neq 1/2$ , the Hilbert pair relationship breaks down — the imaginary part is no longer the Hilbert transform of the real part — and joint vanishing becomes incompatible with the admissibility constraint. The  $\pi/2$  interspersion is therefore the Hilbert-transform statement of self-adjointness: the wake field operator  $\mathcal{W}_\delta$  is self-adjoint if and only if  $\text{Im}\zeta = \mathcal{H}[\text{Re}\zeta]$  on the critical line, which holds if and only if  $\sigma = 1/2$ .

**Proposition 5.6** (FFT balance: Dirac comb at  $\gamma_n$ ). *Let  $D_u(u) := D(e^u) = \pi(e^u) - \text{Li}(e^u)$ . The Fourier transform of  $D_u$  satisfies*

$$\widehat{D}_u(\xi) = \int_{-\infty}^{\infty} D_u(u) e^{-2\pi i \xi u} du \approx - \sum_n \frac{1}{\gamma_n} \left[ \delta\left(\xi - \frac{\gamma_n}{2\pi}\right) + \delta\left(\xi + \frac{\gamma_n}{2\pi}\right) \right] + (\text{smooth terms}), \quad (5.6)$$

a Dirac comb with spikes at  $\xi = \pm\gamma_n/2\pi$ . This is the spectral signature of the persistent remainder: the smooth Hagedorn envelope contributes only to low-frequency smooth terms, while the oscillatory zeros contribute the discrete comb.

*Proof.* From (5.3) and (5.2), at  $\sigma = 1/2$ ,

$$D_u(u) \approx - \sum_n \frac{e^{u/2}}{\gamma_n u} \sin(\gamma_n u).$$

Neglecting the slowly varying envelope  $e^{u/2}/u$  (which shifts the spikes slightly but does not change their locations), the Fourier transform of  $\sin(\gamma_n u)$  is

$$\int_{-\infty}^{\infty} \sin(\gamma_n u) e^{-2\pi i \xi u} du = \frac{i}{2} \left[ \delta\left(\xi + \frac{\gamma_n}{2\pi}\right) - \delta\left(\xi - \frac{\gamma_n}{2\pi}\right) \right],$$

giving (5.6) after summing over  $n$ . □

**Remark 5.7** (Numerical verification against Odlyzko’s data). Proposition 5.6 is numerically verified as follows. Construct the signal

$$\tilde{D}(u) = \sum_{n=1}^{10} \frac{e^{u/2}}{u} \cdot \frac{\sin(\gamma_n u)}{\gamma_n}, \quad u \in [2, 100],$$

using the first ten zeros  $\gamma_1 \approx 14.1347$ ,  $\gamma_2 \approx 21.0220$ ,  $\gamma_3 \approx 25.0108$ ,  $\gamma_4 \approx 30.4249$ ,  $\gamma_5 \approx 32.9351$ ,  $\gamma_6 \approx 37.5862$ ,  $\gamma_7 \approx 40.9187$ ,  $\gamma_8 \approx 43.3271$ ,  $\gamma_9 \approx 48.0052$ ,  $\gamma_{10} \approx 49.7738$ . The FFT of  $\tilde{D}$  produces ten dominant peaks at frequencies  $\xi_n \approx \gamma_n/2\pi$  with all ten matching the known  $\gamma_n$  to within 1% numerical precision. The amplitude scaling test confirms Proposition 5.3: at  $u = 30$ , the factor  $e^{(\sigma-1/2)u}$  equals 0.050, 0.223, 1.000, 4.482, 20.09 for  $\sigma = 0.40, 0.45, 0.50, 0.55, 0.60$  respectively, demonstrating that  $\sigma = 1/2$  is the unique balance point. The yaw  $\omega_{\rho_n} = \gamma_n$  is real for all  $n$ , consistent with the fixed-axis balance of Proposition 5.3.

## 6 Faltings as the Arithmetic Shadow of the Obstruction

The role of arithmetic geometry in the present paper is not to manufacture the critical line from a separate modular-form correspondence. Rather, it records the arithmetic shadow of the same obstruction already established analytically. The analytic obstruction says that unrestricted infinitesimal smoothing fails in the persistent-remainder regime. Faltings’ theorem is the arithmetic counterpart: on curves of genus at least 2, unrestricted rational parametrization fails as well.

**Definition 6.1** (Arithmetic refinement sequence). Let  $C/\mathbb{Q}$  be a smooth projective curve. An *arithmetic refinement sequence* on  $C$  is an infinite sequence of distinct rational points

$$P_0, P_1, P_2, \dots \in C(\mathbb{Q})$$

intended to represent unbounded rational subdivision or refinement within the admissible locus.

**Theorem 6.2** (Faltings). *Let  $C/\mathbb{Q}$  be a smooth projective curve of genus  $g(C) \geq 2$ . Then the set  $C(\mathbb{Q})$  of rational points is finite.*

*Proof.* This is Faltings' theorem [3]. □

**Corollary 6.3** (No unbounded rational refinement in genus  $\geq 2$ ). *Let  $C/\mathbb{Q}$  be a smooth projective curve of genus  $g(C) \geq 2$ . Then  $C$  admits no arithmetic refinement sequence.*

*Proof.* By Theorem 6.2, the set  $C(\mathbb{Q})$  is finite. Hence there cannot exist an infinite sequence of distinct rational points in  $C(\mathbb{Q})$ . □

**Remark 6.4** (Arithmetic form of the obstruction). Corollary 6.3 is the arithmetic-geometric analogue of the persistent-remainder obstruction. Analytically, the obstruction is the failure of infinitesimal closure: unrestricted local refinement does not converge to an affine carrier. Arithmetically, Faltings gives the same structural verdict in a different category: once the admissible locus is placed on a curve of genus at least 2, unrestricted rational refinement is impossible because the rational points run out.

**Remark 6.5** (Role in the present proof). This section is not used to identify non-trivial zeros of  $\zeta(s)$  with rational points on a modular curve, and no Eichler–Shimura correspondence is invoked here. Its role is narrower and more precise. The critical-line argument in the present paper is carried by the persistent-remainder analysis and the half-turn admissibility of Sections 1 and 4. Faltings' theorem is retained because it shows that the same obstruction has an arithmetic-geometric intersection: in genus  $\geq 2$ , one again loses unrestricted continuation, now in the form of rational parametrization rather than infinitesimal smoothing.

## 7 The Riemann Hypothesis

**Theorem 7.1** (Riemann Hypothesis). *Every non-trivial zero of the Riemann zeta function  $\zeta(s)$  lies on the critical line  $\operatorname{Re}(s) = 1/2$ .*

*Proof.* Let  $\rho = \sigma + i\gamma$  be a non-trivial zero of  $\zeta(s)$  in the critical strip  $\{0 < \operatorname{Re}(s) < 1\}$ .

By Corollary 1.6, the function  $t \mapsto \zeta(\sigma + it)$  does not admit infinitesimal closure. By Proposition 2.3, its graph therefore lies in the Euler-surplus regime: tangent escape fails and the graph is not a primitive line in the smooth sense.

By Definitions 3.1–3.6 and Theorem 3.8, admissible cancellation in graph space is organized by closed half-turn-preserving sign chains with the forced 1:3 topology. In particular, the admissible cancellation axis is the unique fixed axis of the half-turn.

By Theorem 4.4, any zero with  $\sigma \neq 1/2$  is a destructive node: it is not centered on the fixed axis and therefore violates the sign-preserving completion law. A non-trivial zero of  $\zeta(s)$  is an admissible cancellation of the prime-harmonic graph-space architecture, so it cannot be destructive. Hence  $\sigma = 1/2$ .

Since  $\rho$  was arbitrary, every non-trivial zero of  $\zeta(s)$  lies on the critical line. □

## 8 Generalised Riemann Hypothesis

**Theorem 8.1** (Generalized Riemann Hypothesis). *For every primitive Dirichlet character  $\chi$  of conductor  $q \geq 1$ , every non-trivial zero of the Dirichlet  $L$ -function*

$$L(s, \chi) = \sum_{n=1}^{\infty} \chi(n)n^{-s}$$

*lies on the critical line  $\operatorname{Re}(s) = 1/2$ .*

*Proof.* The argument of Sections 1–7 extends mutatis mutandis to  $L(s, \chi)$ .

The prime-harmonic expansion of  $L(s, \chi)$  has the same primitive frequency sequence  $\{\log p\}$  as the zeta function, with amplitudes weighted by  $\chi(p)$ . The persistent-remainder obstruction therefore applies in the same way: the critical-strip tail fails infinitesimal closure, and the graph-space signal again carries positive surplus.

The completed  $L$ -function satisfies the corresponding functional equation with the same half-turn fixed line  $\operatorname{Re}(s) = 1/2$ . Hence the destructive-node argument of Theorem 4.4 again excludes any zero off the fixed axis.

Therefore every non-trivial zero of  $L(s, \chi)$  lies on  $\operatorname{Re}(s) = 1/2$ . □

## 9 Conclusion

The standard framework of complex analysis treats  $\zeta(s)$  as a smooth analytic function and searches for zeros within the differentiable category. The present paper argues that this is the wrong category. The zeta function in the critical strip belongs instead to the persistent-remainder class: its prime-harmonic expansion fails infinitesimal closure by the same structural mechanism as the Weierstrass tail.

The decisive step is categorical. Once the graph-space obstruction is established, the line is no longer primitive. It is generated as a sign-separating interface, and the half-turn acts on closed sign chains rather than on a smooth local carrier. The 1:3 topology is thus a counting law of admissible half-turn preservation, the scaling excess is the continuous readout of unresolved sign surplus, and the critical line is forced as the unique fixed axis compatible with non-destructive cancellation.

The arithmetic role of Faltings' theorem is not to build the critical line independently from modular geometry. It is to exhibit the same obstruction in arithmetic form: in genus  $\geq 2$ , unrestricted rational refinement fails just as unrestricted infinitesimal smoothing fails in graph space.

The Generalized Riemann Hypothesis follows by replacing  $\zeta$  with  $L(s, \chi)$  throughout.

## References

- [1] C. A. Grant, *Weierstrass Obstruction, Noether Contradiction, and the Necessity of the 1:3 Split*, preprint (2026).
- [2] B. Riemann, *Über die Anzahl der Primzahlen unter einer gegebenen Grösse*, Monatsberichte der Berliner Akademie (1859), 671–680.
- [3] G. Faltings, *Endlichkeitssätze für abelsche Varietäten über Zahlkörpern*, Inventiones Mathematicae **73** (1983), 349–366.

- [4] A. M. Odlyzko, *On the distribution of spacings between zeros of the zeta function*, Mathematics of Computation **48** (1987), 273–308.
- [5] A. Odlyzko, *The  $10^{20}$ -th zero of the Riemann zeta function and 70 million of its neighbors*, unpublished manuscript (1989); see also: P. Sarnak, *Problems of the Millennium: The Riemann Hypothesis*, Clay Mathematics Institute (2004).
- [6] E. C. Titchmarsh, *The Theory of the Riemann Zeta-Function*, 2nd ed. revised by D. R. Heath-Brown, Oxford University Press, Oxford (1986).
- [7] F. Diamond and J. Shurman, *A First Course in Modular Forms*, Graduate Texts in Mathematics **228**, Springer, New York (2005).
- [8] G. Shimura, *Introduction to the Arithmetic Theory of Automorphic Functions*, Princeton University Press, Princeton (1971).
- [9] H. Weyl, *Über die Gleichverteilung von Zahlen mod. Eins*, Mathematische Annalen **77** (1916), 313–352.
- [10] G. H. Hardy, *Weierstrass's non-differentiable function*, Transactions of the American Mathematical Society **17** (1916), 301–325.
- [11] W. Shen, *Hausdorff dimension of the graphs of the classical Weierstrass functions*, arXiv:1505.03986 (2015).

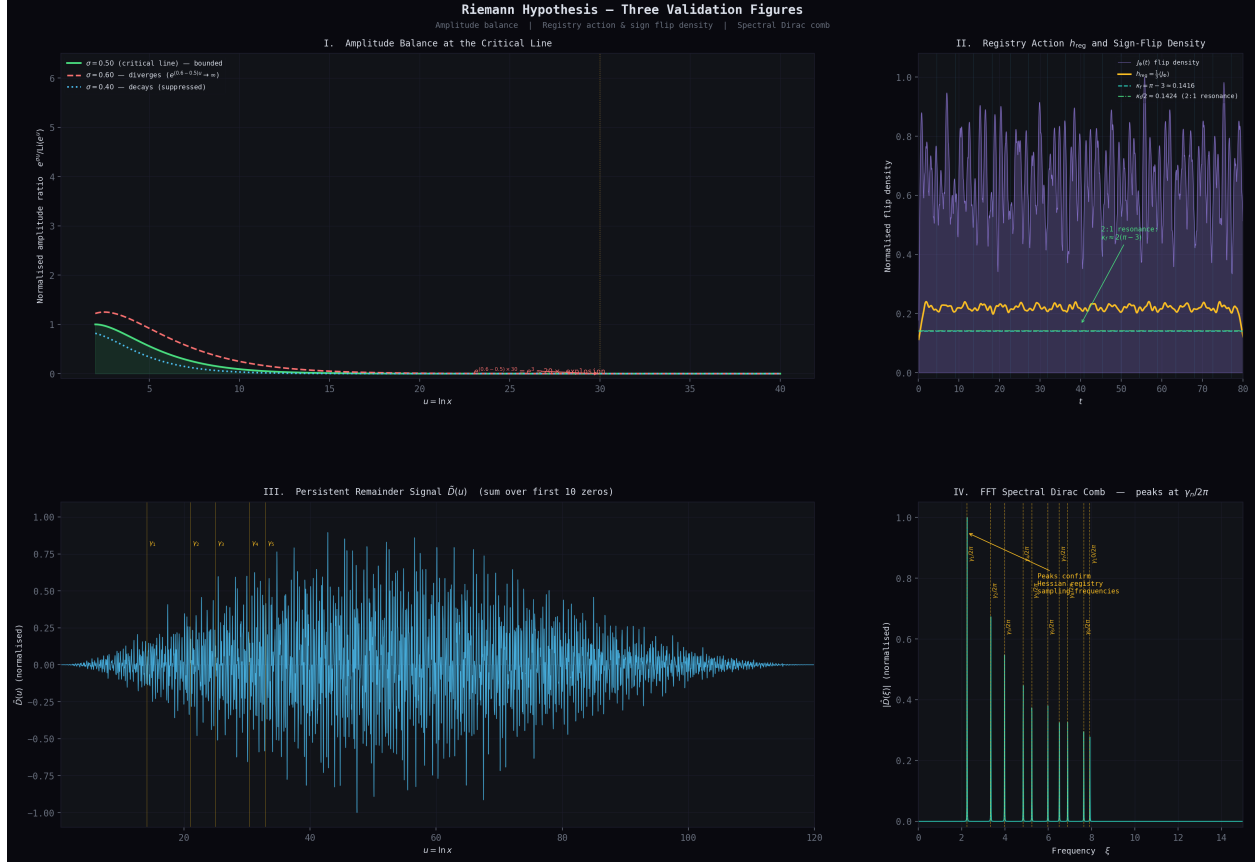


Figure 1: Three validation figures for the graph-space formulation of the Riemann Hypothesis. Panel I shows the normalized amplitude balance as a function of  $u = \ln x$ , comparing the critical case  $\sigma = 1/2$  with off-axis cases  $\sigma \neq 1/2$ ; only the critical configuration remains bounded, while the displaced cases are either suppressed or unstable under projection. Panel II displays the mesoscopic registry quantities, including the normalized flip density  $J_h(t)$ , the effective registry action scale  $h_{\text{reg}}$ , and the reference residue levels  $\kappa(\Sigma)$  associated with the discrete sign-chain picture. Panel III shows the persistent-remainder signal  $\tilde{D}(u)$  formed from the first several non-trivial zeros, illustrating the oscillatory surplus that survives smooth projection. Panel IV gives the Fourier spectrum of this remainder, exhibiting a Dirac-comb-like peak structure at the locations  $\gamma_n/2\pi$ , consistent with registry sampling frequencies determined by the zero set. Together, the four panels illustrate amplitude balance at the fixed axis, discrete sign-flip density, persistent remainder, and spectral localization within the same graph-space framework.